Artificial Neural Network Models for the Evolution of Associative Learning

Emiliano Méndez Salinas*, Franjo Weissing, Magdalena Kozielska
MARM-group, Groningen Institute for Evolutionary Life Sciences, University of Groningen, The Netherlands
*Email contact: e.mendez.salinas@rug.nl

Introduction

Associative learning is the process whereby an organism comes to associate one stimulus or event with other stimuli or events.

The **Rescorla-Wagner Rule** is arguably the most prominent model for explaining how the strength of these associations develop during learning.

Here we use evolving neural networks to address the question **how natural selection shapes associative learning** and whether it will lead to learning patterns that are similar to the Rescorla-Wagner Rule.

Background and Model

Trimmer *et al.* (2012, JTB. 302:39) approached the same question using genetic algorithms and binary trees where learning rules of arbitrary complexity could evolve. We follow their framework but use the more realistic assumption that learning is mediated by a neural network. Their model can be conceptualized as bumblebees that sequentially sample flowers which can either have a nectar reward or not. Each time they experience **reward** (or not) they **update their estimate V of the probability that any given flower provides reward**.

Box 1. The 'learning rule' approach

In the context of Trimmer's model, Rescorla-Wagner updating is given by:

Rescorla-Wagner Rule: $V_{new} = V_{old} + \beta(\lambda - V_{old})$

where λ is the reward (1 or 0) and β is the learning rate. The optimal value of β strongly reflects the number of learning events.

Trimmer et al. (2012) showed that the Rescorla-Wagner Rule readily evolves, even though there is a learning rule with better performance:

Optimal Rule: $V_{new} = V_{old} + \beta(\lambda - 0.5)$

In **our model**, the updating of the probability estimate **V** in response to reward **λ** is not mediated by a learning rule, but by an artificial neural network (see Box 2).

The network has two input nodes (for the reward λ and the previous estimate of V, V_{old}) and one output node (whose value corresponds to the new estimate of V, V_{new}). Information processing happens in-between and is governed by connections between nodes. Connections' weights are genetically encoded and transmitted from parent to offspring (subject to small mutations). Individuals producing a **good estimate** of the true probability of getting nectar have **high fitness** and thus produce more offspring. In this way the population of networks evolves over the generations.

Box 2. The neural network approach

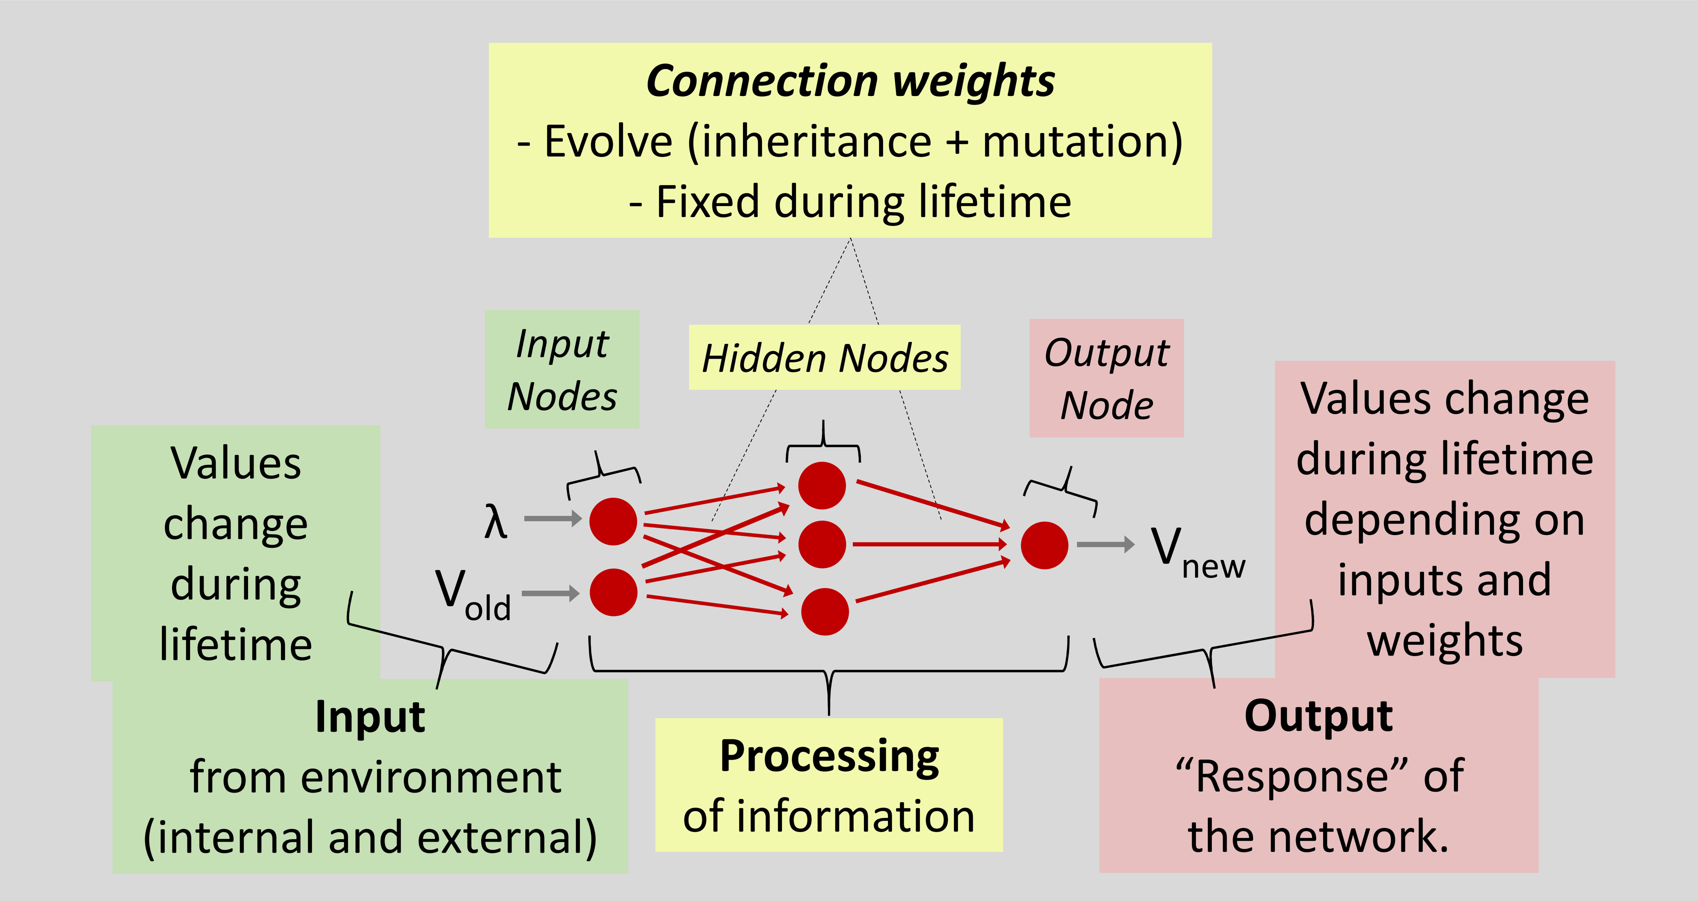

Networks used in our simulations

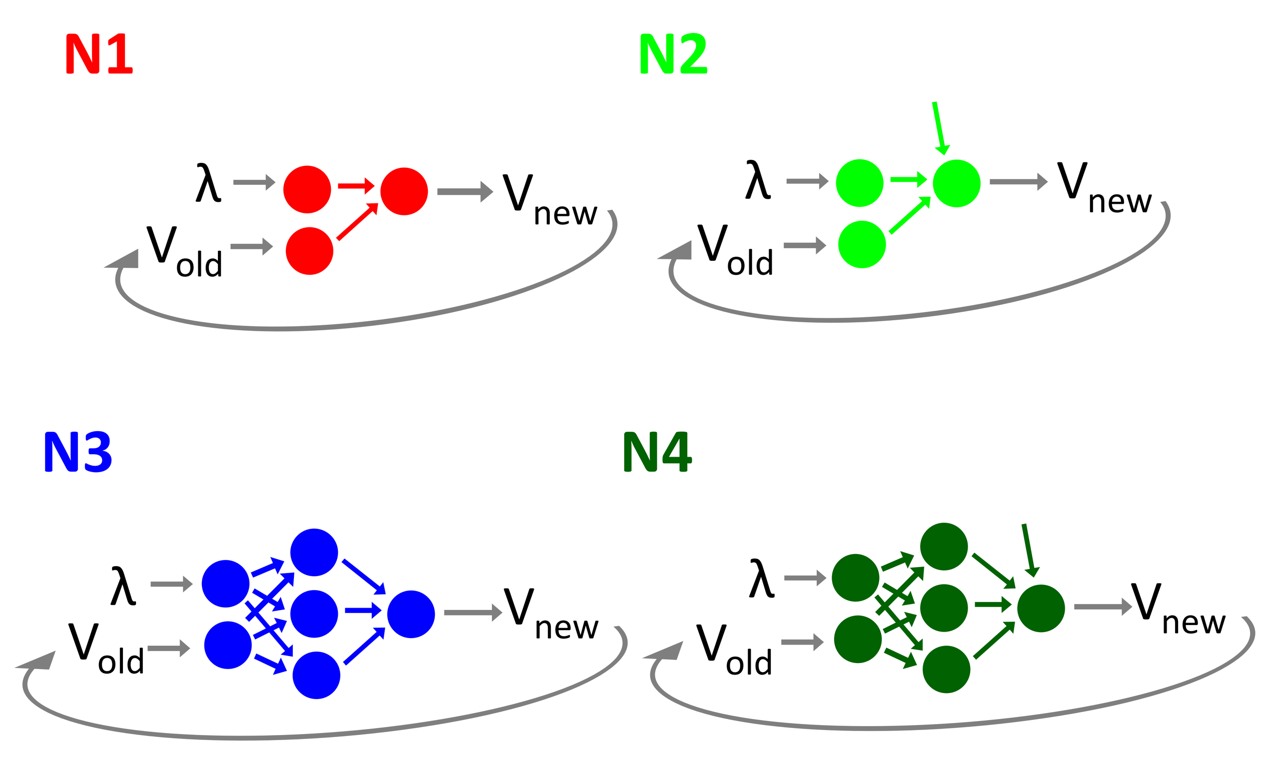

Results

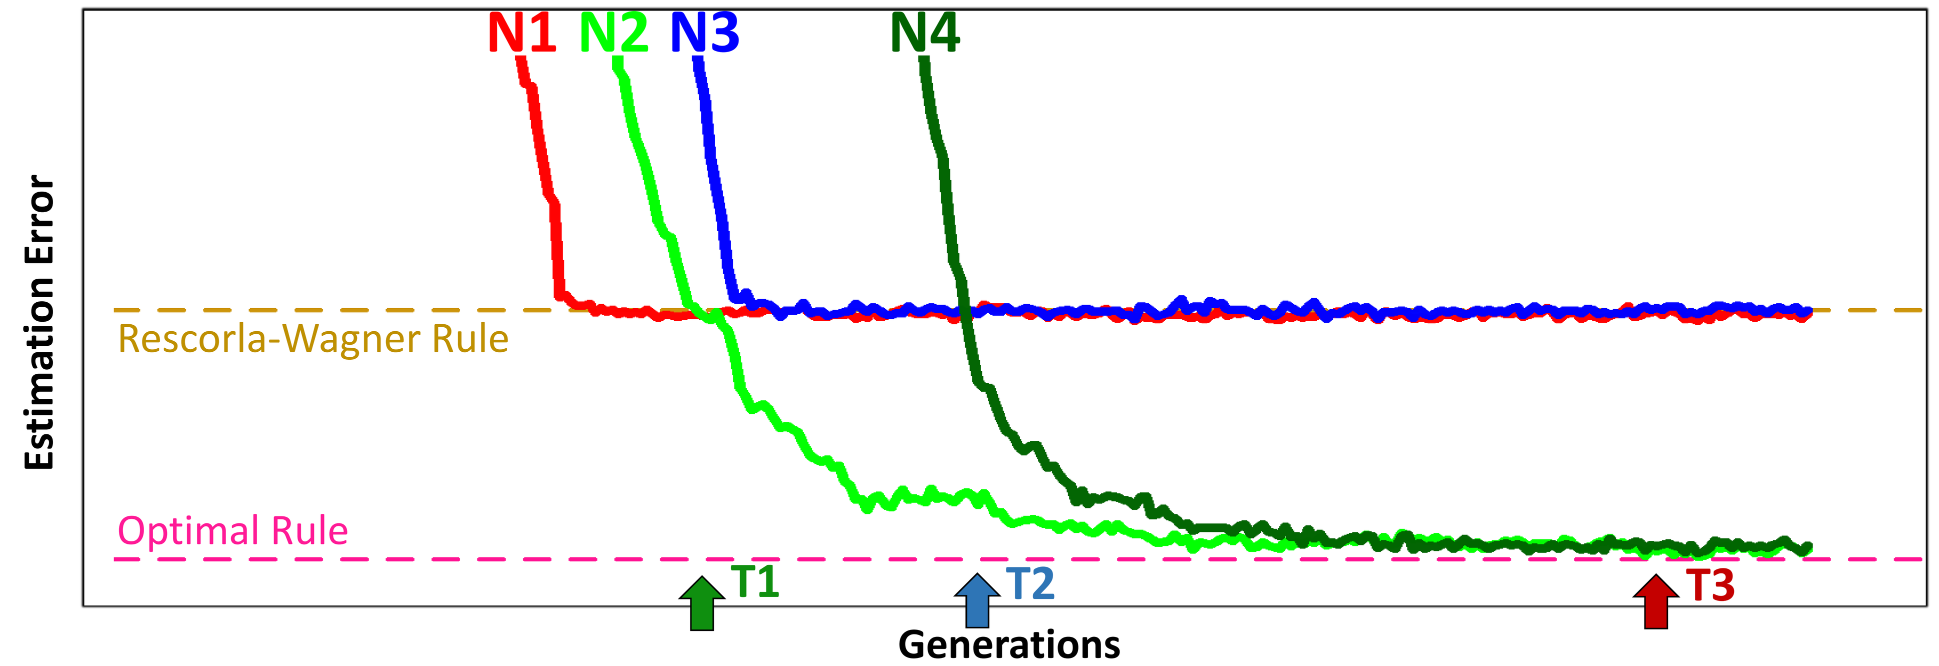

Figure 1. For the networks used, over the generations, the mean estimation error (difference between the true and the estimated value of V) decreases and converges to an asymptotic value.

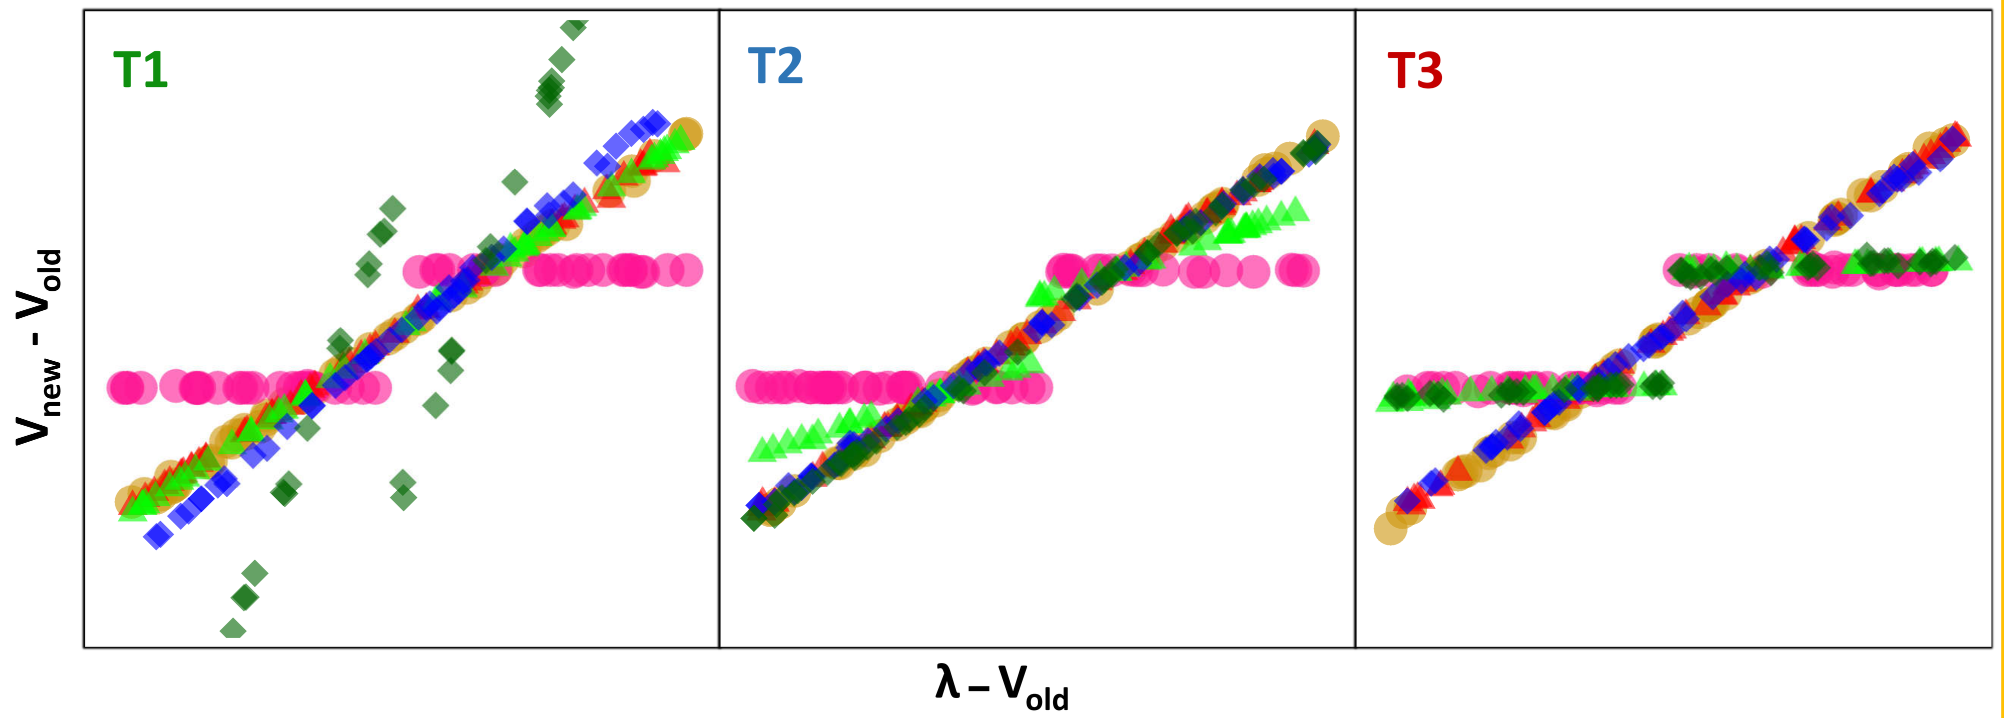

Figure 2. Comparison of the updating mediated by the evolved networks and the Rescorla Wagner Rule. When the difference in estimates V_{new}-V_{old} is plotted against the difference λ-V_{old} between reward and old estimate, the Rescorla-Wagner Rule (gold) produces a straight line with slope β. By plotting the same characteristics in one graph, the updating behaviour of different rules and networks can be compared.

Main Findings and Conclusions

Network **N1** evolves to behave and perform exactly as the **Rescorla-Wagner Rule** (Fig. 1 and Fig. 2).

Network **N2** (**N1** + constant bias) evolves to behave and perform as the **Optimal Rule** (Fig. 1 and Fig. 2 at **T3**).

More complex networks (**N3**, **N4**) evolve more slowly than their simple counterparts (Fig. 1).

But these more complex networks do not perform better (Fig. 1), and they show the same updating behaviour as the simpler networks (Fig. 2 at times **T2** and **T3**).

In line with Trimmer *et al.*'s results, networks that evolve to reach optimal performance, transiently behave and perform as the **Rescorla-Wagner Rule**. (In Fig. 2 **N2** and **N4** at times **T1** and **T2**, respectively).

In a more demanding associative learning task, only some (even more complex) networks outperform the **Rescorla-Wagner Rule** (networks and results not shown).